

**This Page Is Inserted by IFW Operations  
and is not a part of the Official Record**

## **BEST AVAILABLE IMAGES**

**Defective images within this document are accurate representations of the original documents submitted by the applicant.**

**Defects in the images may include (but are not limited to):**

- **BLACK BORDERS**
- **TEXT CUT OFF AT TOP, BOTTOM OR SIDES**
- **FADED TEXT**
- **ILLEGIBLE TEXT**
- **SKEWED/SLANTED IMAGES**
- **COLORED PHOTOS**
- **BLACK OR VERY BLACK AND WHITE DARK PHOTOS**
- **GRAY SCALE DOCUMENTS**

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning documents *will not* correct images,  
please do not report the images to the  
Image Problem Mailbox.**

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau



2

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<b>(51) International Patent Classification <sup>6</sup> :</b> <b>H04L 29/06</b>	<b>A1</b>	<b>(11) International Publication Number:</b> <b>WO 95/17063</b> <b>(43) International Publication Date:</b> 22 June 1995 (22.06.95)
<b>(21) International Application Number:</b> PCT/US94/03979 <b>(22) International Filing Date:</b> 11 April 1994 (11.04.94) <b>(30) Priority Data:</b> 08/169,346 17 December 1993 (17.12.93) US <b>(71) Applicant:</b> TALIGENT, INC. [US/US]; 10201 N. de Anza Boulev. . Cupertino, CA 95014 (US). <b>(72) Inventor:</b> PETTUS, Christopher, E.; 1958 15th Street, San Francisco, CA 94114 (US). <b>(74) Agent:</b> STEPHENS, Keith; Taligent, Inc., 10201 N. de Anza Boulevard, Cupertino, CA 95014 (US).		<b>(81) Designated States:</b> AT, AU, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, ES, FI, GB, HU, JP, KP, KR, KZ, LK, LU, LV, MG, MN, MW, NL, NO, NZ, PL, PT, RO, RU, SD, SE, SK, UA, UZ, VN, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).  <b>Published</b> <i>With international search report.</i>
<b>(54) Title:</b> OBJECT-ORIENTED SECURED COMMUNICATIONS SYSTEM		
<b>(57) Abstract</b>  On a multi-node client server network, a client node obtains access to remote services by means of a communications directory service located in each node of the network. The communications directory service includes a tree structure to which existing directory services and other network services can be added. The tree structure has a plurality of nodes each of which includes specific methods that query and browse the associated directory service if such actions are supported by the underlying service. The communications directory service further includes shared libraries which store a service object associated with each service offered on the network. The service object, in turn, includes the service exchange address and communication link configuration information. A client desiring to access a remote service retrieves the appropriate service object from the communications directory service and uses the service object to set up the communications path.		

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

## OBJECT-ORIENTED SECURED COMMUNICATIONS SYSTEM

### COPYRIGHT NOTIFICATION

Portions of this patent application contain materials that are subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office. All other rights are expressly reserved.

### Field Of The Invention

This invention relates, in general, to distributed computer networks and more specifically to security for distributed network directory and naming services.

### Background Of The Invention

With the tremendous growth of data processing by means of independent, localized data processing devices, such as personal computers and mini computers, data networks have evolved to connect together physically-separated devices and to permit digital communication among the various devices connected to the network.

There are several types of networks, including local area networks (LANs) and wide area networks (WANs). A LAN is a limited area network and data devices connected to a LAN are generally located within the same building. The LAN typically consists of a transmission medium, such as a coaxial cable or a twisted pair which connects together various computers, servers, printers, modems and other digital devices. Each of the devices, which are collectively referred to as "nodes", is connected to the transmission medium at an address which uniquely identifies the node and is used to route data from one node to another. A node which provides resources and services is called a "server" node and a node which uses the resources and services is called a "client" node. A WAN generally encompasses a much larger area and may involve common carrier connections such as telephone lines.

LANs and WANs are often connected together in various configurations to form "enterprise" networks which may span different buildings or locations or extend across an entire continent. Enterprise networks are convenient for several reasons: they allow resource sharing - programs, data and equipment are available to all nodes connected to the network without regard to the physical location of the resource and the user. Enterprise networks may also provide reliability by making several redundant sources of data available. For example, important data files can be replicated on several storage devices so that, if one of the files is unavailable, for example, due to equipment failure, the duplicate files are available.

-2-

One of the most important characteristics of enterprise networks is that they have the capability of bringing a large and sophisticated set of services to all of the attached users for a reasonable cost. However, for the users to exploit the network potential, they must be able to identify, locate and access the network resources.

- 5 When a network is small, locating and accessing the available services is relatively simple, but networks are growing larger and there are many networks that presently very large. Thousand node networks are common and million node networks are on the horizon.

- 10 An example of a very large network is the INTERNET network. which is used by some of the largest public and private organizations. Much of the power of this type of network goes unused simply because the users are either unaware of the facilities available to them or they find the methods of accessing the facilities difficult or confusing. Consequently, in order to assist users in locating and accessing network resources, many existing networks today utilize network directory or naming services
- 15 which accept a resource identifier or name from a user and locate the network address that corresponds to the desired network resource.

- For example, the entered identifier or name can be "descriptive" and specify a resource by describing enough of its attributes to distinguish it from other resources. Such descriptive names are most useful to human users who are searching the
- 20 network for a resource that meets certain specified criteria, but they are also require the most computing resources and are often difficult to distribute effectively. There presently exist a number standards for such descriptive name services. For example, the Consultative Committee on International Telephony and Telegraphy (CCITT) and the International Standards Organization (ISO) have developed a standard for a
- 25 descriptive name service known as X.500

- Naming and directory services (these will be referred to together as "directory services" hereafter) are presently implemented in a variety of ways. The simplest implementation is to use a single, centralized database contained in a local server node to hold a list of names and corresponding network addresses. An example of
- 30 such a localized directory service is shown in Figure 1. Figure 1 illustrates a computer network arranged in a "client-server" configuration comprising a plurality of client nodes 106, 108, 120, 122 and 128 which may, for example, be workstations, personal computers, minicomputers or other computing devices on which run application programs that communicate over various network links including links 102, 110, 116,
- 35 126 and 136 with each other and with server nodes, such as nodes 100, 112, 124, 132 and 138. The server nodes may contain specialized hardware devices and software programs that can provide a service or set of services to all or some of the

-3-

client nodes. The client nodes are the users of the various network services which, in turn, are provided by the server nodes

Typically, the centralized directory service database 104 is located in one of the server nodes, such as node 100. A client node, such as client node 108, can access the directory service by connecting to server node 100, entering a resource identifier or name and retrieving the network address of the associated service. By means of conventional database techniques, a client node may be able to search over the database in order to locate a given resource. In addition, many directory services support browsing by using partial name descriptions, "wild cards" and placeholders. Such centralized directory services with single databases work well in small networks where the number of network addresses is small. However, in larger networks, it is often not feasible to store all the resource identifiers in one central location. Further, a single database represents a single point of failure which can disable the entire network. In addition, a centralized database often suffers from poor performance. For example, while it may be relatively efficient for a local client, such as client 108, to connect to server 100 and access database 104, a remote client, such as client 120, which must link through several servers, 124 and 112, along with a "gateway" link 116, will incur a significant amount of network overhead and the overall system "cost" of the access will be high. With a large number of remote access attempts, directory service provider 100 can quickly become both a processing and communication bottleneck for the entire network.

In order to overcome these problems, additional prior art techniques have been developed which distribute the database data over multiple locations. Such a system is shown schematically in Figure 2. Figure 2 depicts a client-server type of network which is similar to that shown in Figure 1. In particular, elements which correspond in the two figures have corresponding numeral designations. For example, client 108 in Figure 1 is similar to client 208 in Figure 2. The difference between the two networks is that the directory service database has been replicated in a number of the server nodes. For example, server node 200 contains a directory service database 204 as do server nodes 212 (database 214), server node 232 (database 230) and server node 224 (database 218). There are a number of prior art methods for replicating the data in each of the databases. Some systems replicate each resource identifier individually in each database, other systems replicate the entire database. Still other systems replicate individual nodes or limit replication by partitioning the database in some manner.

The distributed system shown in Figure 2 avoids the problems associated with the centralized database. Since the data is replicated, there is no single point of

-4-

failure and, since the data is usually available on a nearby server node, there are no "remote" client nodes and network overhead is greatly reduced.

However, the distributed system has its own problems. For example, some method must be used to insure data consistency if multiple sources can update the databases. Some systems force data consistency by keeping all copies of the data tightly synchronized in a manner similar to a conventional database system. Other system insure data integrity by means of conventional concurrency arbitration schemes.

Such distributed naming and directory services are effective on homogeneous networks in which the same access methods and protocols apply over the entire network. In this case, a consistent set of names and rules can be developed to permit location and access of various resources with relative ease. However, many large networks are heterogeneous - not only do the networks comprise many types of different computers, including work stations, personal computers, mini-computers, super-computers and main frames, but the network itself is often composed of many independent smaller networks which are connected together by interfaces called "gateways". These smaller networks may have their own access methods and protocols. Further, the heterogeneous construction and organization of these large networks does not lend itself to central control and management which could dictate common methods and protocols.

In many large networks which are comprised of a set of smaller networks which are connected together, each of the underlying separate networks may have its own different directory service utilizing a specific protocol. In this type of network a user may have to be familiar with each network directory service protocol and may have to shift from protocol to protocol as searches are performed from network to network. Consequently, in such a heterogeneous network, one of the main difficulties in accessing network resources arises from a lack of a consistent globally-accessible directory of network resources which can operate over heterogeneous networks without involving the user in the details and the protocol involved in accessing each of these separate networks. Today's networking services have various naming and authentication systems. However, there are no unifying apparatus and method to provide support for any security system transparently.

Accordingly, it is an object of the present invention to provide a communication directory security service which provides a single globally accessible directory security service which is capable of interacting with various existing directory security services and other services with existing and future directory security services which are provided on a network.

-5-

### Summary Of The Invention

The foregoing problems are solved and the foregoing objects are achieved in one illustrative embodiment of the invention in which a communications directory security service is located in each node of the network. The communications directory security service includes a tree structure to which existing directory security services and other network services can be added. The tree structure has a plurality of nodes each of which includes specific methods that query and browse the associated directory security service if such actions are supported by the underlying service. The communications directory security service further includes shared libraries which store a service object associated with each service offered on the network. The service object, in turn, includes the service exchange address and communication link configuration information. A client desiring to access a remote service retrieves the appropriate service object from the communications directory security service and uses the service object to set up the communications path.

In one embodiment of the invention, each node uses a reconfigurable protocol stack to establish network connections to remote nodes. The communications directory security service stores a set of stack definitions which allow the reconfigurable stack to be set up for a particular communication link. Each service object corresponding to a particular service, contains reference to one or more stack definitions for communication links appropriate to that service. When a client retrieves the service object, one of the stack definitions is selected based on criteria such as quality of service or availability of the link.

### Brief Description Of The Drawings

The above and further advantages of the invention may be better understood by referring to the following description in conjunction with the accompanying drawings, in which:

Figure 1 is a block schematic diagram of a prior art client server network which incorporates a local directory service.

Figure 2 is a block schematic diagram of a prior art client server network which incorporates a distributed directory server.

Figure 3 is a block schematic diagram of a computer system, for example, a personal computer system in which the inventive object oriented printing interface operates.

Figure 4 is a block schematic diagram of a client server network which incorporates the inventive communications directory service.



-6-

Figure 5 is a detailed block schematic diagram of a prior art protocol stack used to transmit data between two nodes structure in accordance with the International Standards Organization seven layer model.

5 Figure 6 is a block schematic diagram of the major components of the communications directory service.

Figure 7 is a schematic diagram of an illustrative directory tree set up which allows browsing over various directory services and other network services.

Figure 8 is a block schematic diagram of the main components of a server node illustrating how a service program interacts with the communications directory service.

10 Figure 9 is a simplified flowchart of the steps involved in making a new service available on the network.

Figure 10 is an expanded flowchart of the steps carried out by the service program in order to activate a service object.

15 Figure 11 is a block schematic diagram of the main components of a client node illustrating how an application program interacts with the communications directory service to access a service.

Figure 12 is a simplified flowchart of the steps involved in accessing a service available on the network.

20 Figure 13 is an expanded flowchart of the steps carried out by the service program in order to activate a service object.

### Detailed Description Of Illustrative Embodiments

The invention is preferably practiced in the context of an operating system resident on a personal computer such as the IBM PS/2 or Apple Macintosh computer. 25 A representative hardware environment is depicted in Figure 3, which illustrates a typical hardware configuration of a computer 300 in accordance with the subject invention. The computer 300 is controlled by a central processing unit 302, which may be a conventional microprocessor; a number of other units, all interconnected via a system bus 308, are provided to accomplish specific tasks. Although a particular 30 computer may only have some of the units illustrated in Figure 3 or may have additional components not shown, most computers will include at least the units shown.

Specifically, computer 300 shown in Figure 3 includes a random access memory (RAM) 306 for temporary storage of information, a read only memory (ROM) 35 304 for permanent storage of the computer's configuration and basic operating commands and an input/output (I/O) adapter 310 for connecting peripheral devices such as a disk unit 313 and printer 314 to the bus 308, via cables 315 and 312, respectively. A user interface adapter 316 is also provided for connecting input

-7-

devices, such as a keyboard 320, and other known interface devices including mice, speakers and microphones to the bus 308. Visual output is provided by a display adapter 318 which connects the bus 308 to a display device 322 such as a video monitor. The workstation has resident thereon and is controlled and coordinated by operating system software such as the Apple System/7 operating system.

In a preferred embodiment, the invention is implemented in the C++ programming language using object-oriented programming techniques. C++ is a compiled language, that is, programs are written in a human-readable script and this script is then provided to another program called a compiler which generates a machine-readable numeric code that can be loaded into, and directly executed by, a computer. As described below, the C++ language has certain characteristics which allow a software developer to easily use programs written by others while still providing a great deal of control over the reuse of programs to prevent their destruction or improper use. The C++ language is well-known and many articles and texts are available which describe the language in detail. In addition, C++ compilers are commercially available from several vendors including Borland International, Inc. and Microsoft Corporation. Accordingly, for reasons of clarity, the details of the C++ language and the operation of the C++ compiler will not be discussed further in detail herein.

As will be understood by those skilled in the art, Object-Oriented Programming (OOP) techniques involve the definition, creation, use and destruction of "objects". These objects are software entities comprising data elements and routines, or functions, which manipulate the data elements. The data and related functions are treated by the software as an entity and can be created, used and deleted as if they were a single item. Together, the data and functions enable objects to model virtually any real-world entity in terms of its characteristics, which can be represented by the data elements, and its behavior, which can be represented by its data manipulation functions. In this way, objects can model concrete things like people and computers, and they can also model abstract concepts like numbers or geometrical designs.

Objects are defined by creating "classes" which are not objects themselves, but which act as templates that instruct the compiler how to construct the actual object. A class may, for example, specify the number and type of data variables and the steps involved in the functions which manipulate the data. An object is actually created in the program by means of a special function called a constructor which uses the corresponding class definition and additional information, such as arguments provided during object creation, to construct the object. Likewise objects are destroyed by a special function called a destructor. Objects may be used by using their data and invoking their functions.

The principle benefits of object-oriented programming techniques arise out of three basic principles; encapsulation, polymorphism and inheritance. More specifically, objects can be designed to hide, or encapsulate, all, or a portion of, the internal data structure and the internal functions. More particularly, during program design, a program developer can define objects in which all or some of the data variables and all or some of the related functions are considered "private" or for use only by the object itself. Other data or functions can be declared "public" or available for use by other programs. Access to the private variables by other programs can be controlled by defining public functions for an object which access the object's private data. The public functions form a controlled and consistent interface between the private data and the "outside" world. Any attempt to write program code which directly accesses the private variables causes the compiler to generate an error during program compilation which error stops the compilation process and prevents the program from being run.

Polymorphism is a concept which allows objects and functions which have the same overall format, but which work with different data, to function differently in order to produce consistent results. For example, an addition function may be defined as variable A plus variable B ( $A+B$ ) and this same format can be used whether the A and B are numbers, characters or dollars and cents. However, the actual program code which performs the addition may differ widely depending on the type of variables that comprise A and B. Polymorphism allows three separate function definitions to be written, one for each type of variable (numbers, characters and dollars). After the functions have been defined, a program can later refer to the addition function by its common format ( $A+B$ ) and, during compilation, the C++ compiler will determine which of the three functions is actually being used by examining the variable types. The compiler will then substitute the proper function code. Polymorphism allows similar functions which produce analogous results to be "grouped" in the program source code to produce a more logical and clear program flow.

The third principle which underlies object-oriented programming is inheritance, which allows program developers to easily reuse pre-existing programs and to avoid creating software from scratch. The principle of inheritance allows a software developer to declare classes (and the objects which are later created from them) as related. Specifically, classes may be designated as subclasses of other base classes. A subclass "inherits" and has access to all of the public functions of its base classes just as if these function appeared in the subclass. Alternatively, a subclass can override some or all of its inherited functions or may modify some or all of its inherited functions merely by defining a new function with the same form (overriding or modification does not alter the function in the base class, but merely modifies the use

of the function in the subclass). The creation of a new subclass which has some of the functionality (with selective modification) of another class allows software developers to easily customize existing code to meet their particular needs.

Although object-oriented programming offers significant improvements over  
5 other programming concepts, program development still requires significant outlays of time and effort, especially if no pre-existing software programs are available for modification. Consequently, a prior art approach has been to provide a program developer with a set of pre-defined, interconnected classes which create a set of objects and additional miscellaneous routines that are all directed to performing  
10 commonly-encountered tasks in a particular environment. Such pre-defined classes and libraries are typically called "application frameworks" and essentially provide a pre-fabricated structure for a working application.

For example, an application framework for a user interface might provide a set of pre-defined graphic interface objects which create windows, scroll bars, menus, etc.  
15 and provide the support and "default" behavior for these graphic interface objects. Since application frameworks are based on object-oriented techniques, the pre-defined classes can be used as base classes and the built-in default behavior can be inherited by developer-defined subclasses and either modified or overridden to allow developers to extend the framework and create customized solutions in a particular  
20 area of expertise. This object-oriented approach provides a major advantage over traditional programming since the programmer is not changing the original program, but rather extending the capabilities of the original program. In addition, developers are not blindly working through layers of code because the framework provides architectural guidance and modeling and, at the same time, frees the developers to  
25 supply specific actions unique to the problem domain.

There are many kinds of application frameworks available, depending on the level of the system involved and the kind of problem to be solved. The types of frameworks range from high-level application frameworks that assist in developing a user interface, to lower-level frameworks that provide basic system software services  
30 such as communications, printing, file systems support, graphics, etc. Commercial examples of application frameworks include MacApp (Apple), Bedrock (Symantec), OWL (Borland), NeXT Step App Kit (NeXT), and Smalltalk-80 MVC (ParcPlace).

While the application framework approach utilizes all the principles of encapsulation, polymorphism, and inheritance in the object layer, and is a substantial  
35 improvement over other programming techniques, there are difficulties which arise. These difficulties are caused by the fact that it is easy for developers to reuse their own objects, but it is difficult for the developers to use objects generated by other programs. Further, application frameworks generally consist of one or more object "layers" on top

-10-

of a monolithic operating system and even with the flexibility of the object layer, it is still often necessary to directly interact with the underlying operating system by means of awkward procedural calls.

In the same way that an application framework provides the developer with  
5   prefab functionality for an application program, a system framework, such as that  
included in a preferred embodiment, can provide a prefab functionality for system level  
services which developers can modify or override to create customized solutions,  
thereby avoiding the awkward procedural calls necessary with the prior art application  
frameworks programs. For example, consider a printing framework which could  
10   provide the foundation for automated pagination, pre-print processing and page  
composition of printable information generated by an application program. An  
application software developer who needed these capabilities would ordinarily have  
to write specific routines to provide them. To do this with a framework, the developer  
only needs to supply the characteristics and behavior of the finished output, while the  
15   framework provides the actual routines which perform the tasks.

A preferred embodiment takes the concept of frameworks and applies it  
throughout the entire system, including the application and the operating system. For  
the commercial or corporate developer, systems integrator, or OEM, this means all of  
the advantages that have been illustrated for a framework such as MacApp can be  
20   leveraged not only at the application level for such things as text and user interfaces,  
but also at the system level, for services such as printing, graphics, multi-media, file  
systems, I/O, testing, etc.

Figure 4 shows a schematic overview of the illustrative client- server network  
illustrated in Figures 1 and 2 which has now been provided with the inventive  
25   communications directory service which is the subject of the present invention. A  
comparison of Figures 2 and 4 indicates that, although the general network  
configuration is the same for both networks, when the inventive communications  
directory service is used as shown in Figure 4, all of the nodes now include a  
communications directory service (CDS) module. In particular, server nodes 400, 412,  
30   424, 432 and 438 now include the inventive communications directory service  
modules 405, 413, 425, 430 and 439, respectively. In addition, any or all server nodes  
may also include a conventional directory service 404, 414, 418 and 430, respectively.  
These conventional directory services will be referred to as physical directory services  
hereinafter to distinguish them from the communications directory service.

35   In addition to the server nodes, each of the client nodes 406, 408, 420, 422 and  
428 (CDS 407, 409, 421, 423 and 429, respectively) also includes a communications  
directory service. As will hereinafter be explained in detail, in enterprise networks  
such as that shown in Figure 4, information to be sent from one node to another is

-11-

generally divided into discrete messages or "packets" and the packets are transmitted between nodes in accordance with a predefined "protocol". In this context a "protocol" consists of a set of rules or procedures defining how the separate nodes are supposed to interact with each other.

5 In order to reduce design complexity, most networks are organized as a series of "layers" or "levels" so that information passing from one node to another is transmitted from layer to layer. Within each layer, predetermined services or operations are performed and the layers communicate with each other by means of predefined protocols. The purpose for the layered design is to allow a given layer to  
10 offer selected services to other layers by means of a standardized interface while shielding those layers from the details of actual implementation within the layer. As will hereinafter explained in detail, both the client and server nodes cooperate with the communications directory service modules in order to structure the network layers which, in turn, control the type and parameters of the connections between client and  
15 servers in accordance with the type of service being offered.

In an attempt to standardize network architectures (the overall name for the sets of layers and protocols used within a network), a generalized model has been proposed by the International Standards Organization (ISO) as a first step towards international standardization of the various protocols now in use. The model is called  
20 the open systems interconnection (OSI) reference model because it deals with the interconnection of systems that are "open" for communication with other systems. The proposed OSI model has seven layers which are termed (in the order which they interface with each other) the "physical", "data link", "network", "transport", "session", "presentation" and "application" layers. The purpose of the OSI model is to attempt to  
25 standardize the processes conducted within each layer.

In accordance with the OSI model, the processes carried out in the physical layer are concerned with the transmission of raw data bits over a communication channel. The processes carried out in the data link layer manipulate the raw data bit stream and transform it into a data stream that appears free of transmission errors.  
30 The latter task is accomplished by breaking the transmitted data into data frames and transmitting the frames sequentially accompanied with error correcting mechanisms for detecting or correcting errors.

The network layer processes determine how data packets are routed from the data source to the data destination by selecting one of many alternative paths through  
35 the network. The function of the transport layer processes is to accept a data stream from a session layer, split it up into smaller units (if necessary), pass these smaller units to the network layer, and to provide appropriate mechanisms to ensure that the

-12-

units all arrive correctly at the destination, with no sequencing errors, duplicates or missing data.

The session layer processes allow users on different machines to establish "sessions" or "dialogues" between themselves. A session allows ordinary data transport between the communicating nodes, but also provides enhanced services in some applications, such as dialogue control, token management and synchronization. The presentation layer performs certain common functions that are requested sufficiently often to warrant finding a general solution for them, for example, encoding data into a standard format, performing encryption and decryption and other functions. Finally, the application protocol layer contains a variety of protocols that are commonly needed, such as database access, file transfer, etc.

For example, Figure 5 is a diagram of a prior art protocol stack used to connect two nodes in accordance with the OSI standard seven-layer architecture. In accordance with the OSI standard, each protocol stack for a node consists of seven layers. For example, stack 528 for Client node 408 comprises an application layer 500, a presentation layer 502, a session layer 504, a transport layer 506, a network layer 508, a data link layer 510 and a physical layer 512. The operation and the purpose of each of these layers has been previously discussed. Similarly, stack 532 for Server node 432 consists of layers 514-526. Although the actual data communication occurs between the two physical layers 512 and 526 over a data link 530, when the stack is arranged as shown in Figure 5, each layer can be thought of as communicating with its "peer" which is a layer at the same level as a given layer. For example, the application layers 500 and 514 can be thought of as communicating directly even though information passes through all of the layers 502-512, across data link 530 and back through the layers 516-526. Similarly presentation layers 502 and 516 can communicate peer-to-peer.

The protocol stacks such as those shown in Figure 5 are typically implemented using a plurality of data buffers. Each data buffer constitutes a protocol layer in which processing is done. After processing is complete in a layer the data may be transferred to another buffer for further processing in accordance with another layer.

In accordance with one aspect of the invention, the protocol stacks which control peer-to-peer communications in the network are configured by stack definitions stored in the communications directory service. these stack definitions are associated with each service type available on the network so that the protocol stacks can be dynamically configured in response to service requests from the application programs.

A more detailed diagram of a communications directory service module is shown in detail in Figure 6; the module includes three major components: a

-13-

hierarchical directory tree 602 which allows each of the physical directory services and other services provided by the network to be located by means of a conventional tree searching techniques; a set of stack definition objects 604 which are used to program the dynamically reconfigurable stacks that allow a client to request data from a remote server over a prespecified communication link and a set of "service objects" 606. Each service object is associated with one service available on the network and contains the network address or exchange address at which the service is available and a reference 608 to one or more of the stack definition objects 604. As will hereinafter be explained in detail, a reference to one of these service objects is obtained by an application program desiring to access the corresponding service. The information in the object identified by this reference is then sent to the reconfigurable stack in order to set up the communication path.

The stack definitions 604 each consist of a set of layer definitions that specify the processing carried out in each layer and the interactions between the layers. The stack definitions are defined at the time that the communications links are installed on the network system. In particular a stack definition is provided for each different type of communication link on the system.

The communications directory service 600 has a number of other features which make its operation particularly convenient for users on the network. In particular, the directory service stores information regarding the available services and directory services in the directory tree as a series of objects. In addition, the stored libraries in the communications directory service contain objects and, thus, both application programs and service programs can communicate directly with the communications directory service by means of objects. Since the communications directory service is present on each of the nodes, a client need not connect with any directory service server before using the directory service, instead the objects themselves provide the interface.

Further, the use of predefined objects, such as the service objects 606, provides a simple method to allow an application program to open communications with a desired service. The service objects can be sent directly to the networking service in order to open a communications path without requiring the application program to concern itself with the details of the actual path construction. In addition, because any existing physical directory services are encapsulated in the node objects of the directory tree 602, the physical directory services are completely hidden from the client application allowing the client application to interact with the inventive communications directory service with a consistent protocol and feature set.

A more detailed diagram of the directory tree 602 found in the communications directory service is shown in Figure 7. As will hereinafter be explained, the directory



-14-

tree can be traversed by means of simple search commands and full browsing capabilities are provided by using wild cards and placeholders. The directory tree is designed so that objects are returned as the result of searches with the type of object which is returned being determined by the implementation of the portion of the directory which returns the object. Examples of objects which can be returned from a directory tree search are references to the service objects 606 shown in Figure 6 which contain information used to connect to remote service; principle objects which contain security and authentication information about users; "business cards" which contain collections of information about users and other classes which may be added to the directory tree by program developers.

The directory tree is organized as a single hierarchical tree such as that shown in Figure 7. It should be noted that the tree configuration shown in Figure 7 is for an illustrative purposes only and an actual tree configuration can differ significantly from the illustrated configuration without departing from the scope and principles of the present invention. Each of the nodes in the tree is formed by a "namespace" object. A namespace object can refer to other namespace objects or can refer directly to services or other physical directory services. The ultimate root of the directory tree is the root namespace object 700; the root namespace object 700 and the other namespace objects which form the nodes of the tree are inserted into a conventional tree structure which can be traversed to find the node members.

Each node object, such as root namespace object 700, includes methods which facilitate directory searches. In particular, these methods may include a search method which returns an iterator over all entries in the directory, a method which returns entries to which the namespace refers and a method which returns entries that match a given property expression. Additional methods may be provided to obtain the root namespace from lower level directory node.

In the tree directory shown in Figure 7, the ultimate nodes or "leaf" nodes are the physical directory services and the other available network services. As these leaf nodes are also namespace objects, they may include methods which return the associated directory protocol and methods which can interact with the physical directory to search or browse the directory. Since each of the leaf nodes contains methods which are written specifically for the associated service, the methods deal with the protocol issues involved with the associated service allowing the user to interface with the communications directory service in a consistent manner no matter what directory service is involved.

For example, in Figure 7, three nodes 704, 706 and 708 are shown which interact with three separate physical directory services. A fourth node, 710, is also provided, which node is called a "native" namespace node. This latter node contains

-15-

a reference to each of the services that are provided by the network. As will hereinafter be explained, in order to make each service available to the users, it is "registered" in the native namespace 710. "Registration" in the namespace means to insert a reference to the service into the directory where it can found by someone  
5 traversing the directory. Essentially, registration consists of streaming the associated service object into the native namespace where it can be discovered by queries.

In addition, to registering a service, it is also possible to "publish" a service in the directory or in the physical directory services resident in the leaf nodes. Publication differs from registration in that publication involves inserting a limited set of  
10 information about the service into a directory node. This limited set of information may, for example, consist of the service name, type of service and a connection address. In a preferred form of the invention, the native namespace handles publications - when a service entity is registered with the native namespace, the native name space determines in which other namespaces the entity should be published and directs the  
15 publication accordingly.

In order to provide full branching capabilities, intermediate namespaces, such as name space 702 may be inserted between the root namespace 700 and the leaf nodes 704-710. The intermediate namespace refers to other namespaces.

As previously mentioned, the inventive communications directory service  
20 interacts with both service programs and application programs to transparently set up the necessary communication links between the client and server. This interaction involves the creation of service objects in the communications directory service by the service providers and the configuration of the protocol stacks in the server nodes. A client desiring to access a service retrieves the associated service object and uses it to  
25 configure the protocol stacks in the client node to set up the communication link.

Figures 8-10 describe the operations performed by a service program operating in a server node in order to make a new service available on the network for use by application programs running in the client nodes. Figure 8 is a schematic block diagram of the portions of the server node programs that are involved in the creation  
30 and activation of a new service. Figure 9 is an illustrative flow chart which describes the interaction between the service program, the inventive communications directory service and the node networking services in order to establish the new service and to configure the networking service to operate over an appropriate communication link. Figure 10 is a more detailed block diagram illustrating the activation of the new service  
35 by activating the associated service object.

More particularly, Figure 8 illustrates the basic configuration of a server node in providing a service to a remote client. The server node is arranged with a common area, or system address space, 800 which address space would include operating

-16-

system programs and various shared libraries that are used by the service applications running on the system. In particular, the system address space 800 includes the communications directory service 804 and the networking service programs and libraries 816.

5        Also in the server node is the service program which runs in its own address space 810. As will be hereinafter explained, the service program 806 interacts with the communications directory service 804 to create a service object which then resides in the communications directory service 804. This service object is distributed  
10       to all of the other nodes in this system and, thus, is available on a local basis to all of the clients on the network. Because the service object includes the appropriate stack definitions for configuring the networking service 816, the application program together with the communications directory service can set up the communications path using the previously stored definitions without involving a client application program in the construction details.

15       Referring to Figure 9, the creation of a new service begins in step 900 and proceeds to step 902. In step 902, the developer of the service program enables the creation of a new service object containing configuration parameters which are appropriate to the new service. This may be done by subclassing the service object classes located in the directory service. In particular, in order to create a new service  
20       object class, the service program developer specifies a unique name for the service, the type of service and, optionally, various communications link types which can be used to access the service. In accordance with normal object-oriented programming language operation, this subclassing information is included in the service program code during compilation. Therefore, when the service program 806 is installed in the  
25       service program address space 810 in the appropriate server node, the constructor of the service object subclass can be called to construct a service object in the communications directory service 804 as illustrated in step 904 (Figure 9). During the process of calling the constructor, the type and quality of service information is passed to the communications directory service as schematically indicated by arrow 802

30       As previously mentioned, the communications directory service 804 includes a set of stack definitions in shared libraries. These stacks definitions are created when the communications links are defined and are associated with a particular transport mechanism. The stack definitions each consist of a set of layer definitions. The layer definitions control the processing of the data in each layer and the interactions  
35       between layers. Each stack definition completely defines the stack from the transport layer through the physical layer (these layers are described in detail above).

In the process of creating a new service object, the communications directory service 804 uses the type and quality of service information provided by the service

program to construct a session layer and include appropriate references to the stored communications stack definitions as set forth in step 906. If more than one communication link can be used to access the new service, appropriate stack definitions are referenced in the new service object and the session layer is  
5 constructed in order to select one of the communication links based on various criteria, such as the quality of service desired and the availability of the communication link.

Thus, the newly-created service object contains a session layer and the appropriate stack definitions which define the stack from the transport layer to the physical layer. The service object is stored in the communications directory service.  
10 However, at this time, no service address or exchange address is provided in the service object because the service object, although created, is not active. Thus, a service object can be created any point when the service program is installed in the server node, but the service program however does not become activated until further steps are taken as described below.

15 More particularly, in order to activate the new service, the service program instructs the communications directory service 804 to send the stored service object to the dynamically reconfigurable protocol stack (DRPS) 822 in networking service 816 as described in step 908. The manner in which the service object is transferred to the DRPS is illustrated in Figure 8. More particularly, the service program 806 first creates  
20 a service program interface 828 in its own address space 810. The communication between the service program and the service program interface is schematically indicated by arrow 808. The service program interface 828, in turn, creates a configuration data stream 814 which can stream data to a server interface 818 which is permanently available and located in the networking service 816.

25 The service program 806 then retrieves the service object including the network configuration data. The configuration data passes from the communications directory service to the service program interface 828 as indicated by arrow 812. The configuration data is then streamed to the system address space 800 as indicated by arrow 814.

30 Next, as illustrated in step 910, the service program activates the service object which, in turn, causes the service exchange address to be returned to the communications directory service 804. A more detailed description of the activation sequence is illustrated in the flow chart shown in Figure 10. More particularly, the activation sequence starts in step 1000 and proceeds to step 1002. In step 1002, the  
35 communications directory service 804 makes the service available by publishing the service on any underlying physical directory services if this is appropriate. This publication consists of registering the name of the service and, in some cases, the service address in the underlying directory services.

-18-

The routine next proceeds to step 1004 in which the service is registered with the communications directory service. As previously mentioned, such registration involves streaming the service object into the native namespace node of the communications directory service. At this point, a reference to the service is added to the CDS directory tree so that the service can be located by user traversing the tree.

Next, in step 1006, the stack definition references which have been passed to the server interface 818 in networking service 816 by the service program 806 are used to set up the stack definitions that will be used to reconfigure the DRPS 822. In step 1008, the created stack definitions are passed to the DRPS (this operation is shown schematically by arrow 820). At this time the information in the service object is also used to set up the session layer 823.

In step 1010, the exchange address, or the network address, of the session layer 823 in the DRPS 822 is returned to the communications directory service 804. In particular, the networking service 816 obtains the session layer address from the DRPS 822 when the session layer is set up. This address is returned, via the service interface 818, configuration data stream 814, to the service program interface 828. Finally, the exchange address is returned, via the data stream 812, to the communications directory service 804. The exchange address, as previously mentioned, is stored in the associated service object located in the communications directory service 804 so that it can later be retrieved by a client program. At this point, activation is complete and the activation routine shown in Figure 10 ends in step 1012. In step 912, the exchange address is returned to the communications directory service and the service activation routine ends in step 914.

A separate data stream is also set up by the service program at this time so that, at a subsequent time, when a client node requests service from the server node, the incoming service requests can be received. In particular, service request data arriving over the physical communication link 824 is passed through the configured DRPS 822 via a separate request data stream 826 which links the session layer 823 to the service program interface 828. The service program interface 828, in turn, forwards the request, via data stream 808, to the service program 806. Reply information is returned by service program 806, via data stream 808, service program interface 828, request data stream 826, DRPS 822, and physical communication link 824 to the client node.

After a new service has been added to the local communication directory service, the copies on the other nodes must be updated. This updating is carried out in a conventional fashion. For example, it is possible to periodically distribute copies on removable disks to all nodes. However, a more preferred method of distributing copies of the communications directory service would be for the local node which

received the update to broadcast the update information to all other nodes. In order to accomplish this broadcast, the broadcast message includes a special header which causes all of the protocol stacks to be set to a predetermined default protocol. In this manner the broadcast message can be received at all nodes.

5        Figures 11 through 13 illustrate the steps involved when a client application coordinates with the communications directory service to access a remote service. In particular, Figure 11 illustrates the appropriate sections of the client node which are involved in accessing a remote service. As in the server node, the client node includes a system address space 1110 which, in turn, contains the inventive  
10        communications directory service 1112 and a networking service 1118. The application program 1100 runs in its own application address space 1104. The interactions of the application program 1100, the directory service 1112 and the networking service 1118 are described in more detail in the flow charts set forth in Figures 12 and 13.

15        More particularly, the client service access routine starts in step 1200 and proceeds to step 1202. As indicated in step 1202, the client interacts with the communications directory service to get a reference to one of the service objects stored in the communications directory service. This interaction is shown  
20        schematically by arrow 1106 on Figure 11. As previously mentioned, this interaction may involve a user directly if the user searches over the directory tree located in the communications directory service 1112. Alternatively, this interaction may involve an intervening application program which cooperates with the communications directory service to select a service in a visual manner, such as, by dragging a document icon onto a service icon.

25        In any case, in accordance with step 1204, a reference to the service object identified by the communications directory service 1112 is returned to the application program 1100 as shown schematically by arrow 1108. The application program, in turn, creates a client interface object 1126 in preparation for sending the configuration data to the network service 1118. The service object reference is passed by the  
30        communications directory service 1112, via a configuration data stream 1114, to the client interface 1126. From the client interface 1126 the configuration data is streamed over configuration data stream 1116 to a server interface object 1120 located in the networking service 1118. The service interface object 1120 is created when the networking service 1118 is created during system boot up and is permanently resident  
35        in the networking service.

In accordance with step 1206, application program 1100 then activates the service object reference. Figure 13 indicates, in more detail, the steps involved in activating the service object reference. More particularly, the activation routine starts

-20-

in step 1300 and proceeds to step 1302. In step 1302, the service object reference is resolved in any of the underlying physical directory services, if appropriate. This resolution is performed by using the service name located in the service object to search over the underlying directory services and to obtain the network address.

- 5 Alternatively, if the service reference is registered in the native namespace of the communications directory service 1112, then the service address or the service exchange can be obtained directly from the service object reference.

In step 1304, the stack definitions contained in the service object are used by the server interface 1120 and the networking service 1118 to set up protocol stack  
10 layers for configuring the DRPS 1124. Next, in step 1306, the created stack definitions are passed to the DRPS as indicated by arrow 1122. These stack definitions then set up DRPS 1124 and configure the communication link in preparation for sending request and reply data between the application program 1100 and the remote service (not shown in Figure 11).

- 15 Next, in step 1308, the address of the session layer 1123 in the DRPS 1124 is returned to the server interface 1120. The activation routine then finishes in step 1310. Returning to step 1208 of Figure 12, the server interface 1120 exchanges the address of the session layer 1123 for the remote service exchange obtained from the service object reference and returns the remote service exchange (as indicated in step  
20 1208), via configuration data stream 1116, client interface 1126 and data path 1102 to the application program 1100. Thus, when the application program communicates with the remote service, it uses the remote service address passed through the communications directory service to the networking service.

- As set forth in step 1210, a separate data path is set up to send service requests  
25 from application program 1100 to the remote service. This separate data path comprises data path 1102, client interface 1126 and the session layer 1123 of the DRPS 1124. In accordance with step 1212, the request information then sent out over physical communication link 1130 to the remote service location. Reply information returns via DRPS 1124, data stream 1128, client interface 1126 and data path 1102 to  
30 the application program 1100. The service request routine then finishes in step 1214.

### Security

- A server defines a service by creating a TServiceDefinition; a client refers to it using a TServiceReference. In the case where the system is interoperating with  
35 another system, no actual TServiceReference exists. The Communication Directory Service of a preferred embodiment manufactures one instead. In the case of network services, a TNetworkServiceReference object is used to connect to the remote service. This object contains a reference to a subclass of TServiceAdapter; this object is used

-21-

to manufacture the stack definition which will be used to connect to the remote service.

Now, many services require some kind of authentication information, which depends on such things as user names and passwords, to be included as part of a connection process. Right now, it's not clear from the architecture where that information is to be collected, or passed on to the stack definition (in which it, presumably, ultimately needs to reside in order to be communicated to the remote server). A preferred embodiment proposes an extension to the current TNetworkServiceReference / Definition mechanism. This extension provides a callback hook in the form of an object which is invoked at the point of service activation. This object can then do any service-defined task required in order to establish the connection with the remote service.

In particular, TServiceAdapter would be extended to include a new member function: TActivationObject\* TServiceAdapter::CreateActivationObject() const;

15 TActivationObject would have the following protocol:

```
class TActivationObject {
public:
    TClientLayer*      CreateClientLayer(const TStackDefintion&);
};
```

20 The activation object would replace the current thunk that is used to create a new client layer. The base class implementation of TActivationObject::CreateClientLayer would just do the standard, default behavior that TNetworkServiceReference / Definition do already:

```
TActivationObject::CreateClientLayer(const TStackDefintion& sd) {
25     return new(kSameHeap, &sd) TClientLayer(sd);
};
```

However, specialized subclasses could do more interesting things, such as:

```
TNovellNetwareServerActivationObject::CreateClientLayer(const
    TStackDefintion& sd) {
30     // Do we have authentication information already?
    if(!fHaveAuthInfo) {
        // Collect the information from the user.
        TStandardText username, password;
        DoSomeHIVoodoo(username, password);
35
        // Convert it into local authentication information.
        fAuthInfo = MangleInfoIntoNovellCookies(username, password);
    }
}
```



-22-

```
// Create the Novell client layer, inserting the authentication
// information as part of the activation process.

5      TNovellNetwareClientLayer* clientLayer = new(kSameHeap, &sd)
        TNovellNetwareClientLayer(sd, fAuthInfo);

        // Do the authentication dialog, and make sure we passed.
        Boolean wePassedAuth = clientLayer->Authenticate();

10     // If true, we're done. If false, we bail out with an error.
        if(!wePassedAuth) {
            DoSomeHIVoodooAndThrowUpABigDialog();
            return NIL;

15     }

        return clientLayer;

};

20     The foregoing description has been limited to a specific embodiment of this
invention. It will be apparent, however, that variations and modifications may be made
to the invention, with the attainment of some or all of its advantages. Therefore, it is
the object of the appended claims to cover all such variations and modifications as
come within the true spirit and scope of the invention.

25
```

-23-

## CLAIMS

Having thus described our invention, what we claim as new, and desire to secure by Letters Patent is:

- 1 1. A multi-node computer network system for connecting a client node to a server  
2 node over a plurality of different communication links, the computer network  
3 system comprising:
  - 4 (a) a service program located in the server node for offering a service to the client  
5 node;
  - 6 (b) first storage apparatus located in the server node;
  - 7 (c) second storage apparatus located in the client node;
  - 8 (d) communications directory service programs located in the client node and in the  
9 server node, each of the communications directory service programs having a  
10 storage mechanism for storing network configuration information for each  
11 service available on the network in the first and second storage apparatus;
  - 12 (e) apparatus located in the server node for storing a service object in the server  
13 node storage apparatus, the service object comprising a network address for  
14 the server node and a reference to the stored network configuration information;
  - 15 (f) apparatus located in the client node for retrieving a stored service object from  
16 the client node storage apparatus in order to set up a connection between the  
17 client node and the server node using one of the plurality of communication  
18 links; and
  - 19 (g) an authenticator object in the memory containing a protocol to authenticate a  
20 network user.

-24-

- 1 2. A multi-node computer network system according to claim 1 wherein each of  
2 the communications directory service programs further comprises a directory  
3 tree structure having a plurality of node objects and a mechanism responsive to  
4 user-entered criteria for traversing the tree structure to locate selected ones of  
5 the node objects.
- 1 3. A multi-node computer network system according to claim 2, wherein at least  
2 some of the plurality of node objects have a service object stored therein.
- 1 4. A multi-node computer network system according to claim 1, wherein the  
2 server node further comprises a first reconfigurable buffer stack and wherein the  
3 service object comprises configuration information for configuring the first buffer  
4 stack.
- 1 5. A multi-node computer network system according to claim 4, wherein the client  
2 node further comprises a second reconfigurable buffer stack and wherein the  
3 service object comprises configuration information for configuring the second  
4 buffer stack to communicate with the first buffer stack.

-25-

- 1 6. A computer system for use with a multi-node computer network for connecting a  
2 client node to a server node over a plurality of different communication links, the  
3 computer system comprising:
- 4 (a) storage apparatus;
  - 5 (b) a processor;
  - 6 (c) a service program controlling the processor for offering a service to the client  
7 node;
  - 8 (d) a communications directory service program having a directory tree structure  
9 with a plurality of node objects and a mechanism responsive to user-entered  
10 criteria for traversing the tree structure to locate selected ones of the node  
11 objects and a storage mechanism for storing network configuration information  
12 for each service available on the network in the storage apparatus;
  - 13 (e) apparatus for storing a service object in the storage apparatus, the service  
14 object comprising a network address for the server node and a reference to the  
15 stored network configuration information;
  - 16 (f) a reconfigurable protocol buffer stack controlled by the service program in  
17 response to the stored network configuration information to connect to the  
18 network over a selected one of the communication links; and
  - 19 (g) an authenticator object in the memory containing a protocol to authenticate a  
20 network user.

-26-

- 1    7.    A computer system for use with a multi-node computer network for connecting a  
2       client node to a server node over a plurality of different communication links in  
3       response to a user command, the computer system comprising:  
4       (a)   storage apparatus;  
5       (b)   a processor;  
6       (c)   an application program controlling the processor for requesting a service from  
7       the server node;  
8       (d)   a communications directory service program having a directory tree structure  
9       with a plurality of node objects and a mechanism responsive to user-entered  
10      criteria for traversing the tree structure to locate selected ones of the node  
11      objects and a storage mechanism for storing network configuration information  
12      for each service available on the network in the storage apparatus;  
13      (e)   apparatus responsive to the user command for retrieving a service object from  
14      the storage apparatus, the service object comprising a network address for the  
15      server node and a reference to the stored network configuration information;  
16      and  
17      (f)   a reconfigurable protocol buffer stack controlled by the application program in  
18      response to the stored network configuration information to connect to the  
19      network over a selected one of the communication links; and  
20      (g)   an authenticator object in the memory containing a protocol to authenticate a  
21      network user.


- 1 8. A method for connecting a client node to a server node over a plurality of  
2 different communication links in multi-node computer network system, both the  
3 client node and the server node having storage apparatus and a  
4 communications directory service program, each of the communications  
5 directory service programs having a storage mechanism for storing network  
6 configuration information for each service available on the network in the  
7 storage apparatus, the method comprising the steps of:
- 8 A. storing a service object in the server node storage apparatus, the service  
9 object comprising a network address for the server node and a  
10 reference to the stored network configuration information;
- 11 B. retrieving a stored service object from the client node storage apparatus  
12 in order to set up a connection between the client node and the server  
13 node using one of the plurality of communication links; and
- 14 C. authenticating a network user utilizing a protocol stored in an  
15 authenticator object of the memory.

-28-

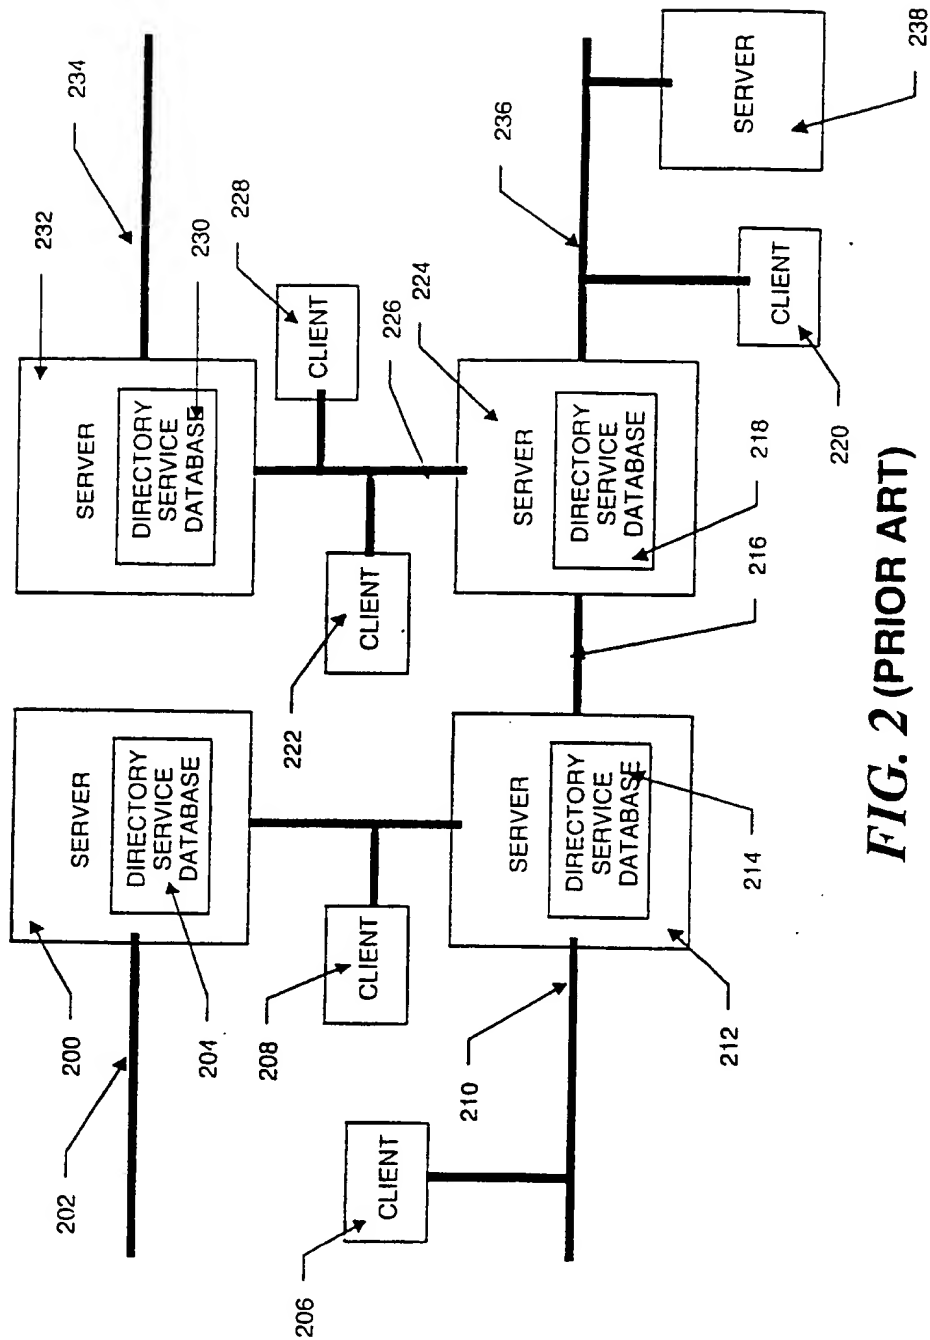
- 1 9. A method for operating a client node in a multi-node computer network for  
2 connecting a client node to a server node over a plurality of different  
3 communication links, the client node having storage apparatus, a processor,  
4 an application program controlling the processor for requesting a service from  
5 the server node and a reconfigurable buffer stack, the method comprising the  
6 steps of:
- 7 A. storing network configuration information for each service available on  
8 the network in the storage apparatus;
  - 9 B. receiving user-entered criteria specifying the service;
  - 10 C. using the user-entered criteria for traversing a directory tree structure  
11 with a plurality of node objects to locate selected ones of the node  
12 objects;
  - 13 D. using the selected node object to retrieve a service object from the  
14 storage apparatus, the service object comprising a network address for  
15 the server node and a reference to the stored network configuration  
16 information;
  - 17 E. configuring the reconfigurable protocol buffer stack to connect to the  
18 network over a selected one of the communication links; and
  - 19 F. authenticating a network user utilizing a protocol stored in an  
20 authenticator object of the memory.

- 1 10. A method for operating a server node in a multi-node computer network for  
2 connecting a client node to a server node over a plurality of different  
3 communication links, the server node having storage apparatus, a processor, a  
4 service program controlling the processor for providing a service with  
5 predetermined characteristics to the client node and a reconfigurable buffer  
6 stack, the method comprising the steps of:
  - 7 A. storing network configuration information for each service available on  
8 the network in the storage apparatus;
  - 9 B. creating a service object from the predetermined characteristics;
  - 10 C. storing the service object in the storage apparatus, the service object  
11 comprising a network address for the server node and a reference to the  
12 stored network configuration information;
  - 13 D. configuring the reconfigurable protocol buffer stack to connect to the  
14 network over a selected one of the communication links; and
  - 15 E. authenticating a network user utilizing a protocol stored in an  
16 authenticator object of the memory.





**FIG. 1 (PRIOR ART)**



**FIG. 2 (PRIOR ART)**

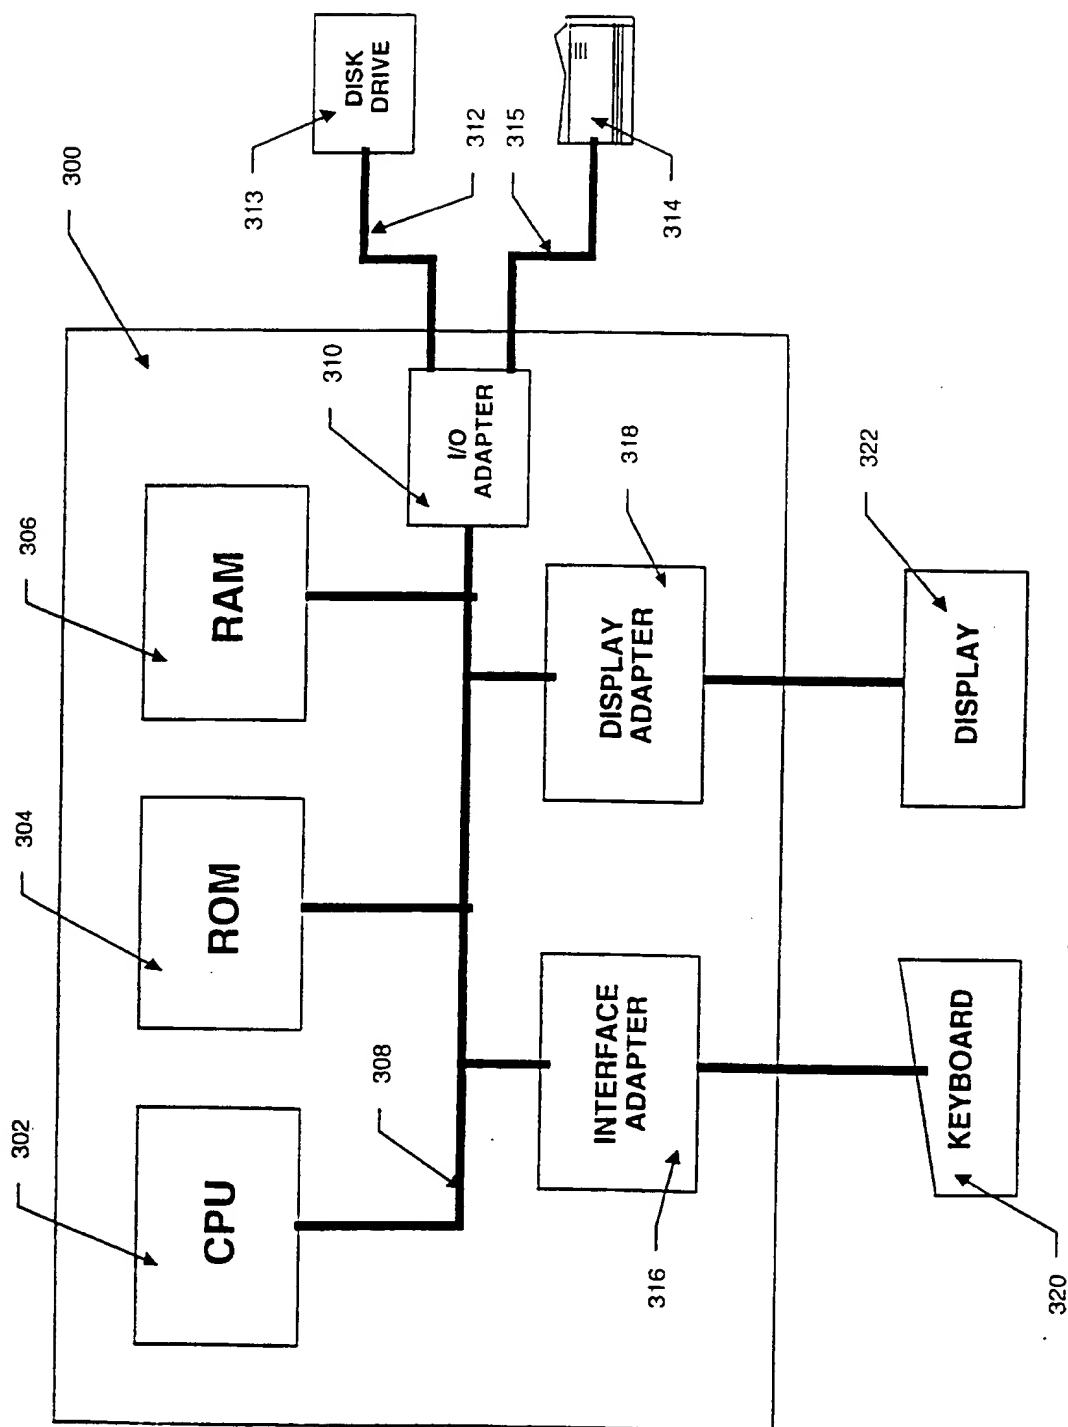





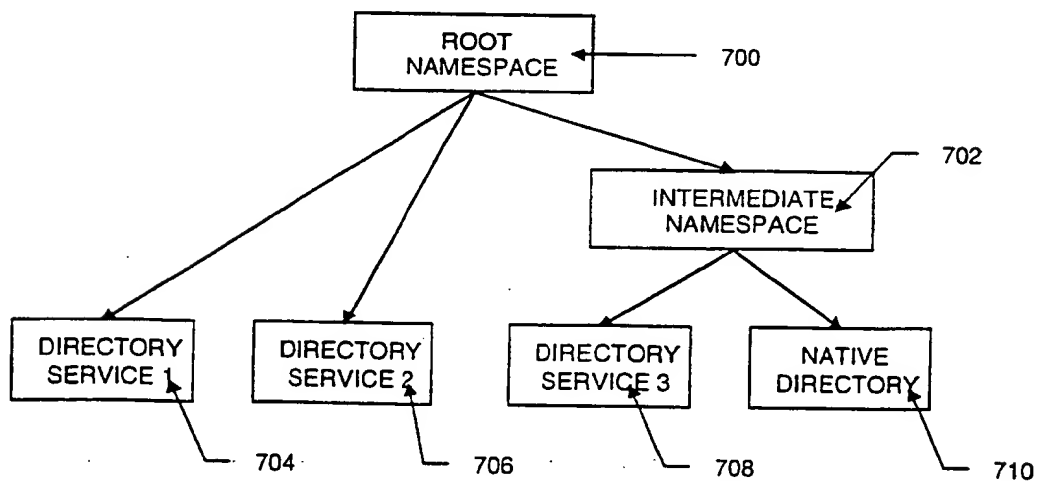
FIG. 3



5/12

**FIG. 5** (Prior Art)

6/12

**FIG. 6****FIG. 7**

7 / 1 2

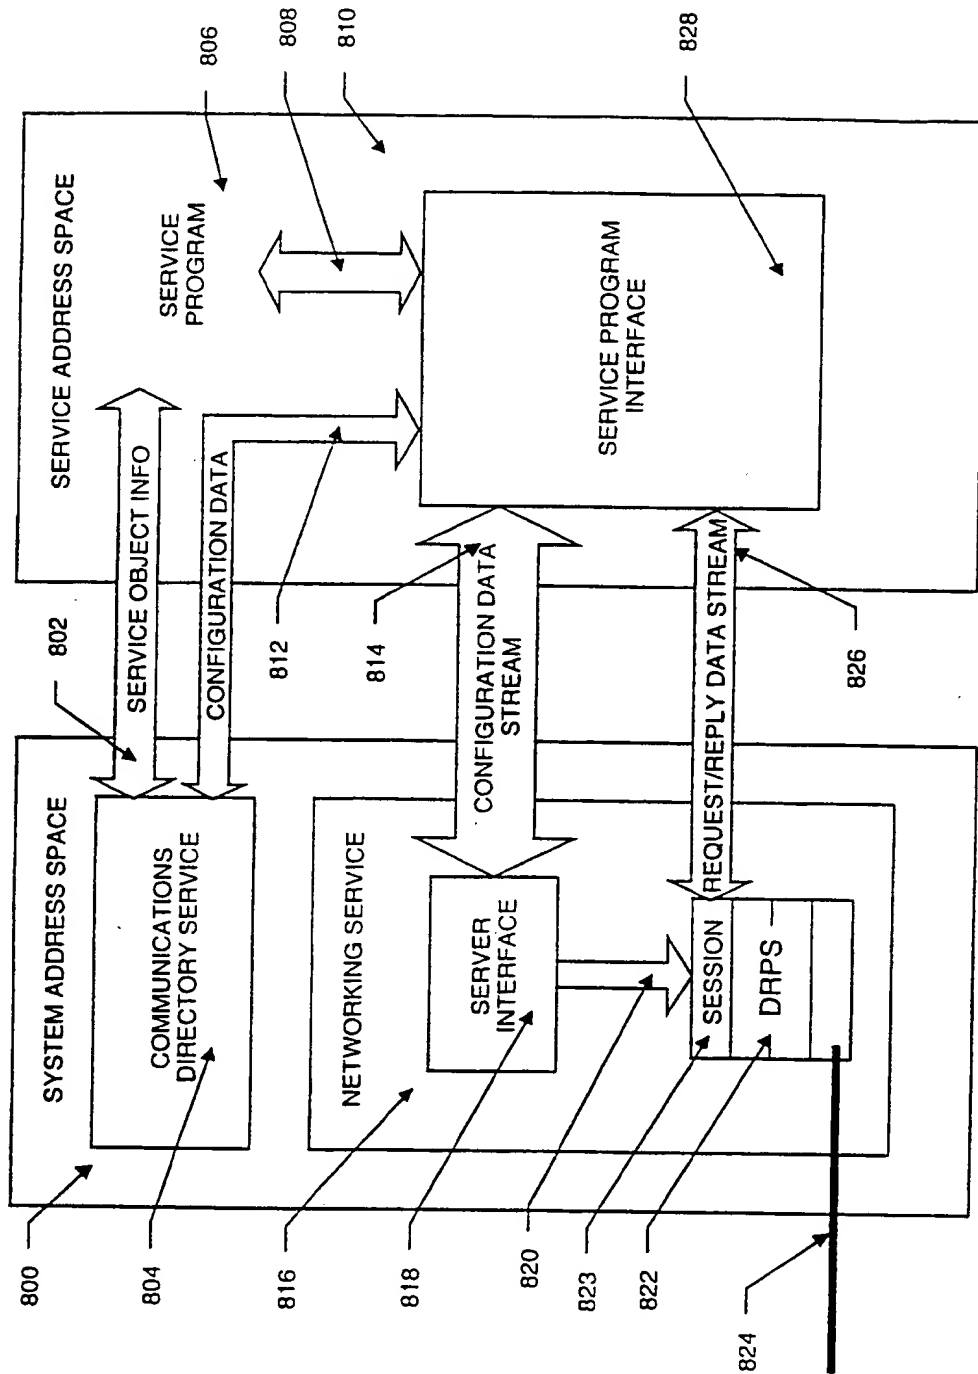
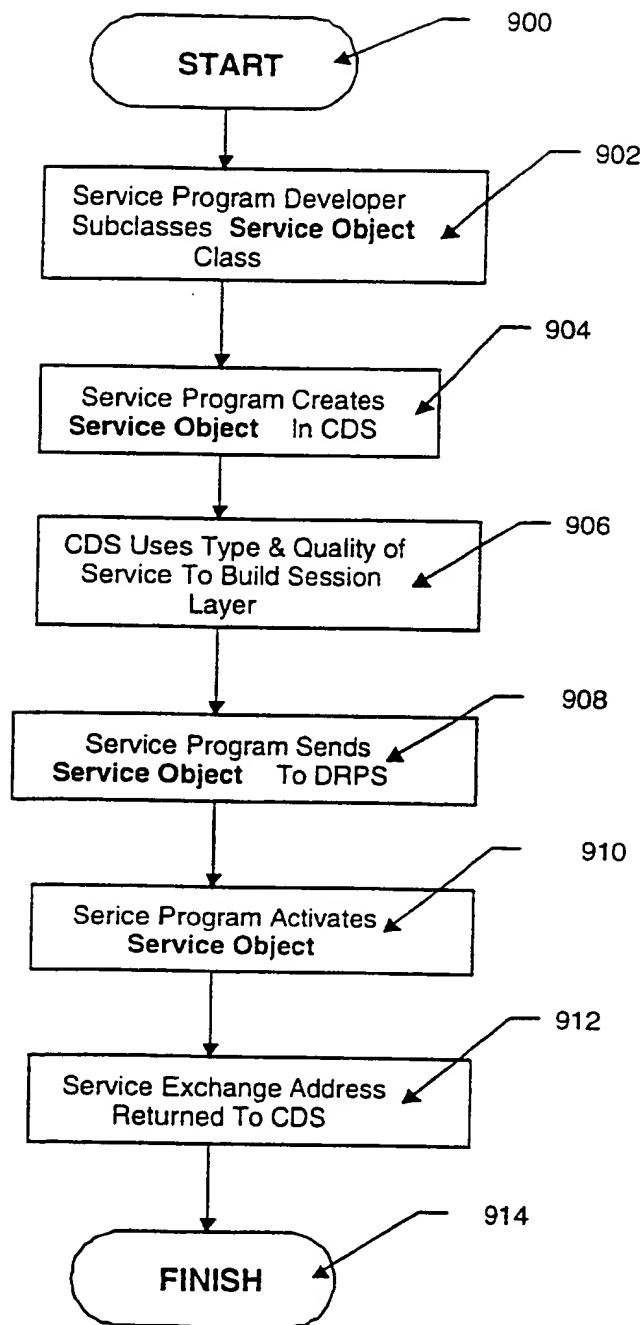
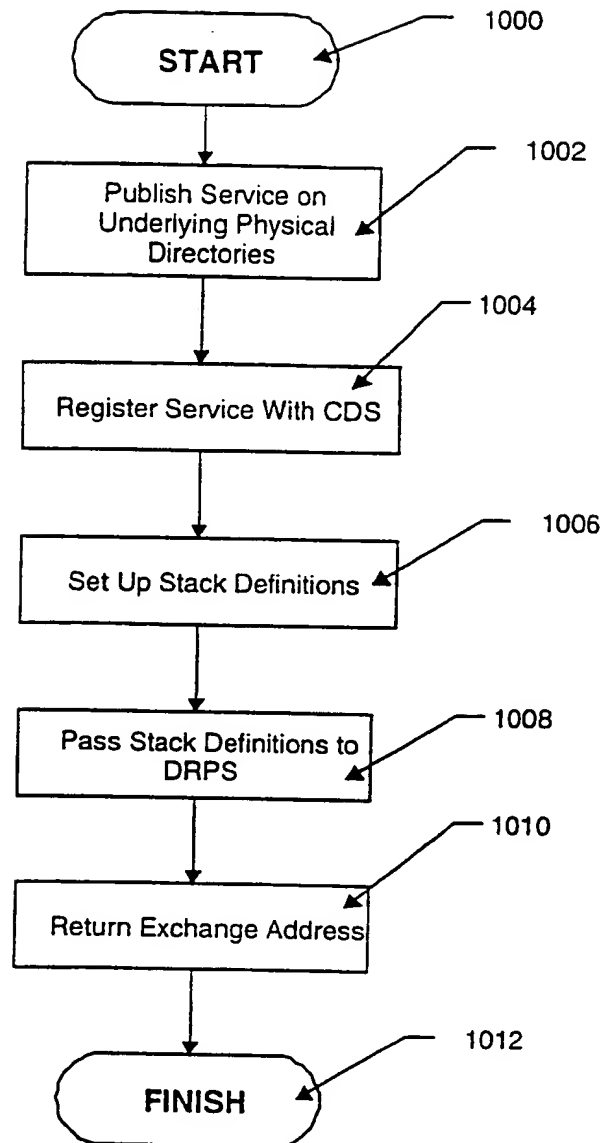


FIG. 8

**8/12****FIG. 9**



**9/12****FIG. 10**

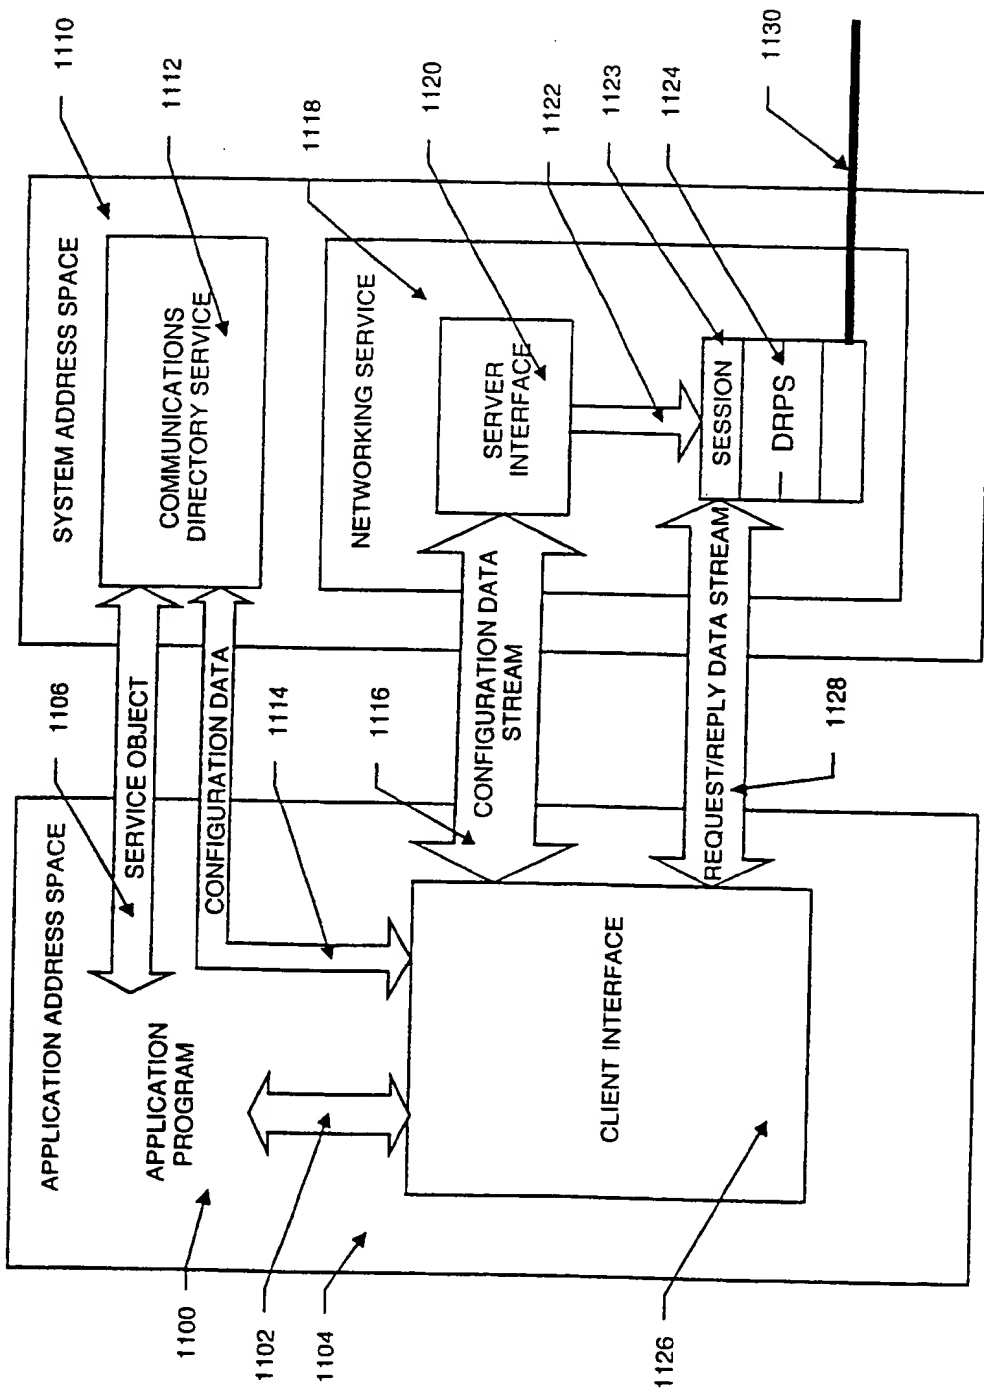
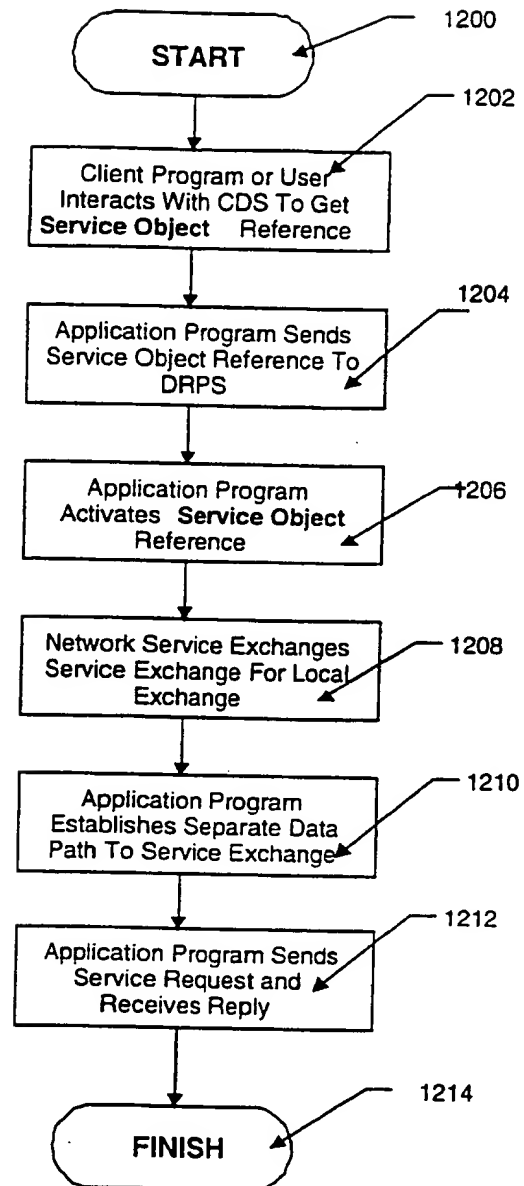
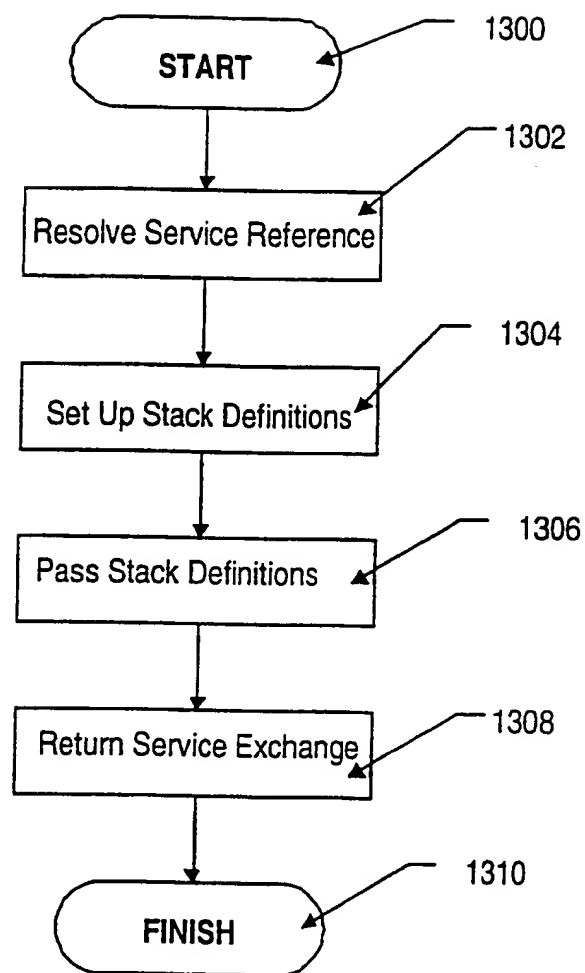


FIG. 11

11/12

**FIG. 12**

12/12

**FIG. 13**

## INTERNATIONAL SEARCH REPORT

Internat. Application No  
PCT/US 94/03979

A. CLASSIFICATION OF SUBJECT MATTER  
IPC 6 H04L29/06

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	COMMUNICATIONS OF THE ASSOCIATION FOR COMPUTING MACHINERY., vol.35, no.7, July 1992, NEW YORK US pages 77 - 98 A.SINHA 'CLIENT-SERVER COMPUTING' see page 88, line 56 - page 89, middle column, line 24 see figure 6	1,6
A	---	2-5,7-10
	--- -/--	

☒ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

## \* Special categories of cited documents:

- \*A\* document defining the general state of the art which is not considered to be of particular relevance
- \*E\* earlier document but published on or after the international filing date
- \*L\* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- \*O\* document referring to an oral disclosure, use, exhibition or other means
- \*P\* document published prior to the international filing date but later than the priority date claimed

- \*T\* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- \*X\* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- \*Y\* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- \*&\* document member of the same patent family

Date of the actual completion of the international search

26 September 1994

Date of mailing of the international search report

06.10.94

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax (+31-70) 340-3016

Authorized officer

Canosa Arete, C

# INTERNATIONAL SEARCH REPORT

Inter-  
national Application No  
PCT/US 94/03979

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	'BLUE BOOK, VOLUME VIII, FASCICLE VIII.8, RECOMMENDATION X.509, THE DIRECTORY - AUTHENTICATION FRAMEWORK' 25 November 1988 , CCITT, THE INTERNATIONAL TELEGRAPH AND TELEPHONE CONSULTATIVE COMMITTEE , GENEVA, CH see paragraph 5 see paragraph 9	1,6
A	-----	2-5,7-10